

# Mapping and Pseudo-Inverse Algorithms for Data Assimilation

Paul Fieguth  
Systems Design Engineering, University of Waterloo

Dimitris Menemenlis and Ichiro Fukumori  
Ocean Science Research Element, Jet Propulsion Laboratory

## I. Overview

### Abstract:

Among existing ocean data assimilation methodologies, reduced-state Kalman filters are a widely-studied compromise between resolution, and computational feasibility. Such reduced-state filters require mapping operators from the fine grid to the reduced state and vice-versa; that is, that the state-reduction and interpolation operators be pseudo-inverses of each other. This poster investigates a variety of approaches to computing the pseudoinverse and also evaluates the mapping performance of eleven interpolation kernels.

### Introduction:

Goal: to understand and predict the general circulation of the oceans.

Existing approaches remain a compromise between resolution, optimality, error specification, and computational feasibility. Widely-studied compromise: reduced-state Kalman filter in which the measurement update takes place on a reduced state compared to the full state of the Ocean General Circulation Model (OGCM).

Main challenge: require mapping operators from the fine (OGCM) state to the reduced state and vice-versa. Let  $\mathbf{x}_f$  and  $\mathbf{x}_c$  represent the fine and coarse state vectors. State reduction  $\mathbf{B}^*$  and interpolation  $\mathbf{B}$  operations defined such that

$$\mathbf{x}_c = \mathbf{B}^* \mathbf{x}_f, \quad \mathbf{x}_f = \mathbf{B} \mathbf{x}_c + \epsilon, \quad \mathbf{B}^* \mathbf{B} = \mathbf{I} \quad (1)$$

$\mathbf{B}^*$  and  $\mathbf{B}$  are pseudoinverses, a condition which ensures that repeated subsampling and interpolation do not lead to a degradation of the coarse-scale data:

$$\mathbf{x}_c = \mathbf{B}^* \mathbf{B} \mathbf{x}_c \quad (2)$$

Objective: to define fast, storage-efficient methods of finding  $\mathbf{B}^*$  from  $\mathbf{B}$ .

Existing mapping and pseudo-inverse schemes often involve the brute-force computation:

$$\mathbf{B} = \mathbf{B}^* \mathbf{T} (\mathbf{B}^* \mathbf{B}^* \mathbf{T})^{-1}, \quad \mathbf{B}^* = (\mathbf{B}^* \mathbf{T} \mathbf{B})^{-1} \mathbf{B}^* \mathbf{T} \quad (3)$$

Where the matrices are of size  $n_f \times n_c$ , where  $n_f$  and  $n_c$  are the fine-grid dimension of the ocean model and the coarse-grid dimension of the reduced state, respectively.

Magnitude of challenge: Suppose we have a global problem with  $1/12^\circ$ -spacing:  $n_f \approx 10^7$ . Suppose the coarse grid has grid spacing of  $2^\circ$ :  $n_c \approx 10^4$ . Then the mapping and pseudo-inverse operations, stored as dense matrices, are each 1 TERABYTE in size!

### Inversion Criteria:

In addition to a computationally efficient approach to identifying a pseudoinverse, the interpolation kernel in  $\mathbf{B}$  must satisfy at least two other requirements.

First: sensitivity to lateral translations must be minimized, to ensure that a slow, advective flow is not progressively corrupted by repeated mapping-interpolations:

$$\mathbf{S} \mathbf{B} \mathbf{B}^* \approx \mathbf{B} \mathbf{B}^* \mathbf{S}, \quad (4)$$

where  $\mathbf{S}$  represents a spatial translation on the fine scale. This is effectively an antialiasing or bandlimiting criterion.

Second: insensitivity to noise, that is, we wish to limit the coarse-scale amplification of fine-scale perturbations. The noise sensitivity is proportional to

$$\frac{|\mathbf{x}_c - \mathbf{x}_c|}{|\delta|} \frac{|\mathbf{x}_f|}{|\mathbf{x}_c|} = \frac{|\mathbf{B}^* \delta|}{|\delta|} \frac{|\mathbf{B} \mathbf{x}_c|}{|\mathbf{x}_c|} \quad (5)$$

The upper bound for this sensitivity is given by the condition number of  $\mathbf{B}$  or  $\mathbf{B}^*$ :

$$\sigma_{\max}(\mathbf{B}) * \sigma_{\max}(\mathbf{B}^*) = \text{cond}(\mathbf{B}) = \text{cond}(\mathbf{B}^*) \geq 1. \quad (6)$$

## II. Fast Inversion

### FFT:

Computing the pseudoinverse by brute force requires enormous storage and computational effort. A simple intuitive approach is to use the FFT:

$$\mathbf{x}_f = \mathcal{F}_2^{-1}[\mathcal{W}(k_x, k_y) \mathcal{F}_2(\uparrow \mathbf{x}_c)], \quad (7)$$

$$\mathbf{x}_c = \downarrow \mathcal{F}_2^{-1}[\mathcal{W}^*(k_x, k_y) \mathcal{F}_2(\mathbf{x}_f)]. \quad (8)$$

Very efficient and fast, however it makes strict stationarity and periodicity assumptions, are incompatible with irregularities (e.g., coastlines).

### Subsampling:

Subsampling methods allow a straightforward alternative to the brute-force approach; define  $\mathbf{x}_s$  of intermediate resolution:

$$\mathbf{x}_c \xleftarrow{\mathbf{B}_0} \mathbf{x}_s \xleftarrow{\mathbf{B}_1} \mathbf{x}_f \quad (9)$$

$$\mathbf{x}_c \xrightarrow{\mathbf{B}_0} \mathbf{x}_s \xrightarrow{\mathbf{B}_1} \mathbf{x}_f$$

Key Idea — The pseudoinverse of  $\mathbf{B}_1^*$  is very easily found:

$$\mathbf{B}_1 = \mathbf{B}_1^{*T} + \mathbf{B}_2, \quad (10)$$

such that a row in  $\mathbf{B}_2$  is zero if the corresponding row of  $\mathbf{B}_1^{*T}$  is non-zero. Problem: the subsampling operator introduces aliasing and leads to substantial shift-sensitivities.

### Implicit Inversion:

Implicit methods avoid explicitly computing  $\mathbf{B}^*$  from  $\mathbf{B}$ , i.e.,

$$\mathbf{x}_c = \mathbf{B}^* \mathbf{x}_f = (\mathbf{B}^* \mathbf{B})^{-1} \mathbf{B}^* \mathbf{x}_f = \mathbf{Q}^{-1} (\mathbf{B}^* \mathbf{x}_f) \quad (11)$$

However even the “small” dense matrix  $\mathbf{Q}^{-1}$  can be unwieldy, both for storage and inversion complexity, for global-sized problems.

### Iterative Inversion:

Instead, we propose to iteratively solve the linear system

$$\mathbf{Q} \mathbf{x}_f = \mathbf{x}_c \quad (12)$$

which is vastly simpler because of the sparsity of  $\mathbf{Q}$ . We apply the Conjugate Gradient method because of its efficiency and simplicity.

Following table compares storage and computational complexity for  $100 \times 100$  coarse-scale and  $1000 \times 1000$  fine-scale problem:

	Storage $\mathbf{B}^*, \mathbf{Q}^{-1}, \mathbf{Q}$	Initialization Effort	Effort Per Mapping
Brute Force	$n_c \cdot n_f$ 100 GB	$n_c^2 \cdot n_f^2$ $10^{13}$	$n_c \cdot n_f$ $10^{10}$
Implicit Method	$n_c^2$ 1 GB	$n_c^3$ $10^{12}$	$n_c^2 + \alpha^2 n_f$ $10^8$
Iterative Method	$\alpha^2 n_c$ 1 MB	$\alpha^3 n_f / \beta$ $10^7$	$\alpha^2 n_f + \alpha^2 n_c$ $2 \times 10^7$

The iterative approach offers tremendous reduction in storage and computational complexity!

Actual reduction in complexity depends on sparsity of  $\mathbf{Q}$  and  $i$ , the number of conjugate-gradient iterations required for convergence:

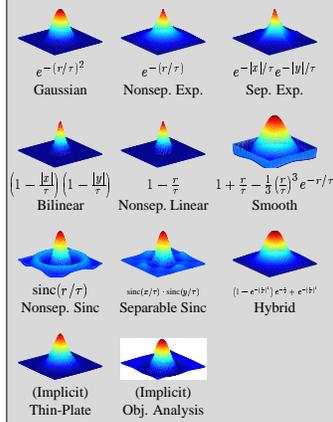
Problem Size	$\mathbf{Q}$ Density	Interpolator Size $\tau$ (fine-scale pixels)							
		2	3	5	8	12	17	28	
$33 \times 33$	0.09	4	6	11	41	174	303	240	
$29 \times 29$	0.12	3	6	11	43	165	291	245	
$25 \times 25$	0.15	3	6	11	41	169	283	233	
$21 \times 21$	0.21	3	6	11	40	158	290	223	
$17 \times 17$	0.30	3	6	11	41	155	238	195	
$13 \times 13$	0.45	3	6	11	38	115	232	168	
$9 \times 9$	0.73	4	6	11	27	117	172	115	

We show the average number of conjugate-gradient iterations to achieve a root-mean-squared accuracy of 0.5%.

## III. Kernels

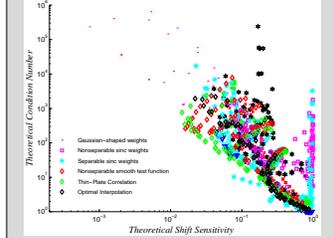
### Kernels Tested:

We have evaluated the shift and noise sensitivities for eleven different interpolation kernels:



### Kernel Assessments:

All tests were carried out in  $20 \times 20$ -coarse-scale,  $200 \times 200$ -fine-scale domains. The theoretical tests measure aliasing (4) and condition number (6):



We can validate these tests experimentally. The shift sensitivity is defined as the root-mean-square ratio

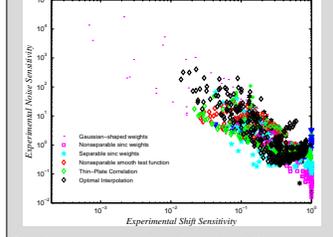
$$\frac{\text{rms} \{ \mathcal{S}(\mathbf{B} \delta_i) - \mathbf{B} \mathbf{B}^* \mathcal{S}(\mathbf{B} \delta_i) \}}{\text{rms} \{ \mathbf{B} \delta_i \}} \quad (13)$$

where  $\delta_i$  is a coarse unit-vector with pixel  $i$  set to one and the rest to zero.

Noise sensitivity is measured by computing the reaction to noise:

$$\frac{\text{rms}(\mathbf{B}^* \mathbf{N}_f)}{\{\sum_i w^2(i)\}^{1/2}}, \quad (14)$$

where  $\mathbf{N}_f$  is an array of unit-variance, independent, Gaussian random variables.

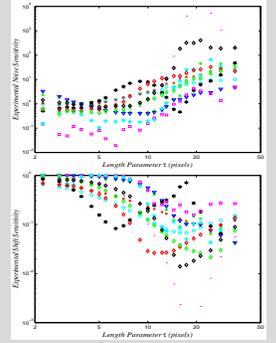


Surprisingly, common kernels such as bilinear, exponential, Gaussian, and sinc functions performed only moderately well.

## IV. Results

### Scale Sensitivity:

A summary illustration of the sensitivity of various interpolants to the choice of scale. Generally, a larger scale leads to smoother interpolants, less aliasing (shift sensitivity), and larger condition number:

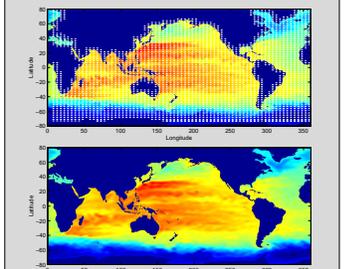


### Kernel Conclusions:

Based on our test results we propose that the Hybrid, Thin-Plate, or Objective Analysis kernels have superior properties and should be recommended for mapping exercises:

Weight	Positivity	Properties	Comments
Gaussian	+	+	Numeric issues
Nonsep. Exp.	+	-	
Separable Exp.	+	-	
Bilinear	-	-	
Cone-shaped	-	-	
Neg-lobe	-	-	
Nonsep. Sinc	-	-	
Sep. Sinc	+	+	Regular Grid
Smooth	+	+	Recommended
Thin-Plate	+	+	Recommended
Optimal Interp.	+	+	Recommended

### Real Data Example:



Mapping test for global-scale problem. We have a  $71 \times 62$  coarse grid and a  $2160 \times 960$  fine grid. The centered locations of the 3551 interpolants are shown as white dots in the top panel; each interpolant has a footprint of  $121 \times 81$  pixels, or  $20 \times 13$  degrees. The bottom panel shows the result of fine-coarse-fine mapping.